

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 January 2002 (17.01.2002)

PCT

(10) International Publication Number
WO 02/05116 A2

55223

U.S. PTO
10/756874
011404

(51) International Patent Classification⁷: **G06F 17/00**

(21) International Application Number: PCT/US01/21075

(22) International Filing Date: 2 July 2001 (02.07.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/611,099 6 July 2000 (06.07.2000) US

(71) Applicant (for all designated States except US): **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; Telefonplan, S-126 25 Stockholm (SE).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **KUFTEDJIAN, Ohaness** [US/US]; 2 Dorchester West, Irvine, CA 92620 (US). **CHENG, Wen, Chung** [—/US]; 7662 El Rio Verde, La Palma, CA 90623 (US).

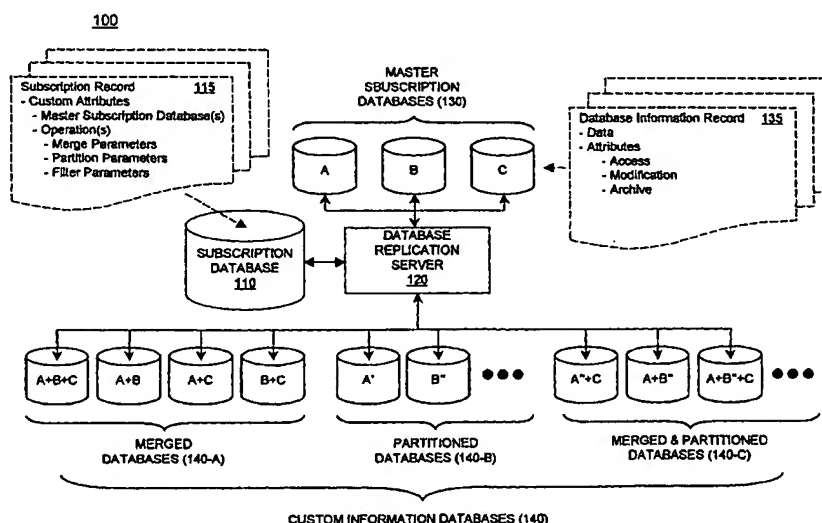
Published:

— without international search report and to be republished upon receipt of that report

(74) Agents: **BURLEIGH, Roger, S.** et al.; Ericsson Inc., 6300 Legacy, MS EVW 2-C-2, Plano, TX 75024 (US).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: AUTOMATIC DATABASE REPLICATION SERVER AND METHODS



(57) Abstract: Systems and methods for automatically replicating database information. A subscription database (110) is queried by a database replication server (120) to obtain custom attributes defined in a plurality of custom database information subscriptions (115). The custom attributes for each database information subscription (115) include: (a) the identification of one or more master subscription databases (130), and (b) at least one operation to be performed on the one or more master subscription database to create a custom information database. The operation(s) to be performed can include merging database information contained in records (135) from two or more master subscription databases (130), partitioning database information contained in one or more master subscription databases (130). The operations identified by the custom attributes for each custom database information subscription (115) are used to automatically generate custom information databases (140) containing preferred database information from the master subscription databases (130).

WO 02/05116 A2

AUTOMATIC DATABASE REPLICATION SERVER AND METHODS

TECHNICAL FIELD OF THE INVENTION

The present invention is directed, in general, to information storage
5 systems, and, more specifically, to systems and methods for automatically
replicating database information.

BACKGROUND OF THE INVENTION

The Information Revolution is now at the point at which the Industrial
10 Revolution was in the early 1820s, about forty years after James Watt's improved
steam engine (first installed in 1776) was first applied, in 1785, to an industrial
operation – the spinning of cotton. And the steam engine was to the first
Industrial Revolution what the computer has been to the Information Revolution
– its trigger, but above all its symbol. Almost everybody today believes that
15 nothing in economic history has ever moved as fast as, or had a greater impact
than, the Information Revolution. But although the computer has provided the
motive power behind the Information Revolution, it is the cargo moved through
the power of computers, *information*, which truly provides benefits.

Much of the valuable information available through the Internet, or any
20 network of computers, is stored in databases. Such databases may store, for
example, telephone directory information, such as a telephone subscriber's
name, address, and telephone number. Typically, separate databases are used
to store directory information for different calling areas, based on area codes or
metropolitan areas. Thus, if a subscriber's calling area is unknown, a search for
25 a directory listing for that subscriber must be independently performed on each
database to locate the listing. In addition, because such calling area databases
may be physically remote, processing delays may be encountered in establishing
and retrieving directory information from remote database servers.

In addition to telephone directory information, databases are used to store
30 other types of information, such as customer purchasing history or other
information relating to an individual's or business' commercial activities.

Databases can also store data related to personal interests in various topical information, such as political, financial, or sports news. In addition to these various databases often being physically remote, a particular database most often only contains a limited amount of information related to an individual or
5 business; *i.e.*, the whole of the information available for a particular individual or business may be distributed among multiple remote databases.

Although the Information Revolution has brought about the general availability of vast amounts of information, the tools for finding and managing such information are still evolving. The tools available to search through
10 tremendous amounts of information often require submitting multiple queries to many different databases. Accordingly, there is a need in the art for systems and methods to combine information from multiple databases according to a user's specific requirements. Preferably, such systems and methods should provide for the automatic creation of databases customized to meet a user's particular
15 parameters.

SUMMARY OF THE INVENTION

To address the above-discussed deficiencies of the prior art, the present invention relates to systems and methods for automatically replicating database
20 information. In the exemplary embodiments described in detail hereinafter, a subscription database is queried to obtain custom attributes associated with each of a plurality of custom database information subscriptions. The custom attributes for each database information subscription include a) the identification of one or more master subscription databases, and b) at least one operation to be
25 performed on the one or more master subscription databases to create a custom information database. The operation(s) to be performed can include merging database information from two or more master subscription databases, partitioning database information contained in one or more master subscription databases, and filtering database information from one or more master
30 subscription databases. The operations identified by the custom attributes for each custom database information subscription are used to generate custom

information databases containing preferred database information stored in the master subscription databases.

The merging, partitioning and filtering operations specified in the custom attributes can be defined on a record or record field basis. For example, 5 complete records from two or more master subscription databases can be merged or, alternatively, one or more fields from records in a first master subscription database can be merged with other fields in records in a second master subscription database.

In one exemplary embodiment, generation of a custom information 10 database is performed by i) creating a local database for storing the custom database information; ii) establishing a connection to each of the master subscription databases identified by the custom attributes; iii) retrieving and processing master database records from each of the master subscription databases to yield custom database records according to the operation(s) 15 identified by the custom attributes; and iv) storing the custom database records in the local database. In some embodiments, the custom database information stored in the local database is published to a remote database; a remote database can be, for example, a database stored on a system at the custom database information subscriber's physical location.

20 An exemplary database replication server for automatically replicating database information, according to the principles of the present invention, is also described hereinafter. A database information server can be a general or specific purpose computing system operative to perform the software-definable processes disclosed herein. An exemplary database information server includes 25 a processor, volatile memory, and a non-volatile media for storing a software-definable process. The processor and volatile memory are operative to retrieve and execute the software-definable process to query a subscription database to obtain custom attributes associated with each of a plurality of custom directory information subscriptions. The operations identified by the custom attributes, 30 described hereinabove, for each custom directory information subscription are

executed by the processor to generate custom information databases containing preferred directory information stored in the master subscription databases.

The foregoing has outlined, rather broadly, the principles of the present invention so that those skilled in the art may better understand the detailed description of the exemplary embodiments that follow. Those skilled in the art should appreciate that they can readily use the disclosed conception and exemplary embodiments as a basis for designing or modifying other structures and methods for carrying out the same purposes of the present invention. Those skilled in the art should also realize that such equivalent constructions do not depart from the spirit and scope of the invention in its broadest form, as defined by the claims provided hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, reference is now made to the following detailed description taken in conjunction with the accompanying drawings, in which:

FIGURE 1 illustrates an exemplary system for automatic database replication to generate a custom information database according to the principles of the present invention;

FIGURE 2 illustrates an exemplary application of the principles of the present invention to automatically generate a custom information database;

FIGURES 3-a and 3-b illustrate an exemplary method for automatically replicating database information to generate a custom information database according to the principles of the present invention; and

FIGURE 4 illustrates a block schematic of an exemplary database replication server and illustrative custom attribute operations for automatically replicating database information to generate a custom information database according to the principles of the present invention.

DETAILED DESCRIPTION

Database replication is the process of creating and managing duplicate versions of a database. A replication process not only copies a database, but can also synchronizes a set of replicas so that changes made to one replica are reflected in all the others. The beauty of replication is that it enables many users to work with their own local copy of a database but have the database updated as if they were working on a single, centralized database. For database applications where users are geographically widely distributed, replication is often the most efficient method of database access. Although database replication can solve many inefficiency problems associated with remote database access, a single database may not contain all of the information desired by a user. For this reason, it may be necessary for a user to manually access multiple databases to retrieve and manage the information required in a particular situation – which presents another form of inefficiency. The present invention is directed to solving the inefficiencies associated with remote database access, as well as providing an automatic means of generating a database containing desired information from many different remote databases.

Referring to FIGURE 1, illustrated is an exemplary system 100 for automatic database information replication according to the principles of the present invention. The core of the system is embodied in a Database Replication Server 120 and a Subscription Database 110. The Database Replication Server 120 can be a general, or specific, purpose computing system operative to perform the software-definable processes described hereinafter; the architecture and operation of general and specific purpose computers are known to those skilled in the art and, thus, a detailed description herein is not necessary to an understanding of the present invention. In general, a Database Information Server 120 includes a processor, volatile memory, such as random access memory (RAM), and a non-volatile media, such as a magnetic storage device, for storing a software-definable process. The processor and volatile memory are operative to retrieve and execute the software-definable process.

The Database Replication Server 120 has access to one or more Master Subscription Databases 130, such as databases A, B and C. The Master Subscription Databases 130 can be locally connected to the Database Replication Server 120 or, more typically, are stored on one or more remote systems accessible via a network, such as the Internet. The data stored in the Master Subscription Databases 130 can be of any type; in particular, the principles of the present invention are advantageously applied to databases containing directory information. As used herein, a "directory" is broadly defined to include a collection of alphabetical or classified lists, such as a listing of telephone subscriber names and their associated telephone numbers and addresses. A "directory" can also be a listing of customers, and can include, for example, a history of each of customer's previous purchases. Within a database, information is logically organized into records, and each record contains predefined fields that store related information. For example, a record can be associated with a telephone subscriber, and contain an address field and a telephone number field. A database, or each record within such database, can also have associated attributes that define, or limit, access to, or modification of, the database or individual records therein.

According to the principles of the present invention, the Database Replication Server 120 also has access to a Subscription Database 110. The Subscription Database 110 can be queried by the Database Replication Server 120 to obtain custom attributes associated with each of a plurality of custom database information subscriptions; the custom attributes for each subscription are stored in a database record (or "subscription record") 115. The custom attributes for each custom database information subscription include: a) the identification of one or more Master Subscription Databases 130, and b) at least one operation to be performed on the one or more Master Subscription Databases 130 to create a custom information database 140. The operation(s) to be performed can include merging database information from two or more Master Subscription Databases 130, partitioning database information contained in one or more Master Subscription Databases 130, and filtering database

information from one or more Master Subscription Databases 130. The operations identified by the custom attributes for each custom database information subscription are used to generate custom information databases 140 containing preferred database information stored in the Master Subscription
5 Databases 130.

As illustrated in FIGURE 1, a subscription can contain custom attributes that instruct the Database Replication Server 120 to automatically generate one or more Custom Information Databases 140 containing merged database information (140-A) from two or more Master Subscription Databases (*i.e.*,
10 A+B+C, A+B, A+C, and B+C). In addition, the custom attributes can instruct the Database Replication Server 120 to automatically partition one or more of the Master Subscription Databases 130. Partitioning of databases provides multiple independent replicas (140-B) of the Master Subscription Databases 130, and allows for different users to simultaneously access and manipulate the database
15 information, while providing both security and privacy between such users. The custom attributes can also specify a filter to be applied to the database records in a Master Subscription Database 130; a filter can be used, for example, to instruct the Database Replication Server 120 to replicate only a subset of the records in a Master Subscription Database 130, such as records in which a zip
20 code field matches a predefined value.

A particularly advantageous feature of the present invention is the generation of Custom Information Databases 140-C containing both merged and partitioned database information from Master Subscription Databases 130; *e.g.*, A'+C is a merged directory containing records from a partitioned replica of
25 Master Subscription Database A and records from Master Subscription Database C. The possible combinations, and advantages, of merged / partitioned directories are essentially limitless.

Reference is now made to FIGURE 2, which illustrates an exemplary application of the principles of the present invention to automatically generate a
30 Custom Information Database 240. Database Replication Server 220 accesses a Subscription Database 110 (see FIGURE 1) to obtain custom attributes

associated with a custom database information subscription; the custom attributes for the subscription are stored in a database record (or "subscription record") 215. In the example illustrated in FIGURE 2, a Master Subscription Database 230-A contains records that include fields for the name and purchase history information (*i.e.*, "Last," "Total," and "Average" purchase values) for customers. Exemplary Master Subscription Database 230-B contains records that include fields for name, address, and area code (AC) and telephone number; *i.e.*, the general information contained in telephone directories.

According to the example in FIGURE 2, the subscriber associated with Custom Information Database 240 desires to create a database that includes records for each customer in Master Subscription Database 230-A that has made purchases totaling greater than \$500; the subscriber also desires to include the address, AC and telephone number for each of those customers. The desired information is not available in any single database, and thus the Custom Information Database 240 must be created from multiple Master Subscription Databases (*i.e.*, Master Subscription Databases 230-A and 230-B).

In order to create the Custom Information Database 240, the exemplary custom attributes in Subscription Record 215 include an instruction for Database Replication Server 220 to filter Master Subscription Database 230-A for records in which a customer's Total Purchases exceeds \$500. The result of this operation can be stored in an intermediate database A_1 ; as used herein, an "intermediate" database is a database created to hold database records that are intended to undergo further operations during the creation of a Custom Information Database. In the example provided in FIGURE 2, the intermediate database A_1 would include the records 1 and 2 from Master Subscription Database 230-A. The exemplary custom attributes in Subscription Record 215 also include an instruction for Database Replication Server 220 to partition Master Subscription Database 230-B; the result of this operation is intermediate database B' . Finally, the exemplary custom attributes in Subscription Record 215 include an instruction for Database Replication Server 220 to merge the address, AC and telephone number fields from each record in intermediate database B' for which

there is a corresponding record in intermediate database A_1 ; *i.e.*, add the address, AC and telephone number fields to each record in intermediate database A_1 to create Custom Information Database 240. As illustrated, the resulting Custom Information Database 240 comprises the desired records of all customers having a purchase total greater than \$500, including the address, AC and telephone number for each of those customers.

Although illustrated in FIGURE 2 as a series of textual instructions, the custom attributes identified in a Subscription Record can be provided in many different formats; a selected format is a function of the specific implementation of the associated Database Replication Server. First, the custom attributes can be identified by preselected codes, or instruction codes, that can be used as input to predefined software functions provided by the Database Replication Server. For example, an instruction code "MERGE (X, Y)" could be interpreted by the Database Replication Server to partition a Master Subscription Database Y, and then merge with a Master Subscription Database X. Similarly, a custom attribute "FILTER(X, Average, >500)" could be used as an instruction code to direct a Database Replication Server to filter all records from Master Subscription Database X in which an "Average" field has a value greater than 500. As an alternative to using codes, or instruction codes, as custom attributes that are used as input to predefined software functions provided by a Database Replication Server, the custom attributes identified in a Subscription Record can be in a direct form of software instructions executable by the Database Replication Server, or in a form of scripting language that can be interpreted by the Database Replication Server to perform the desired operations on each designated Master Subscription Database. Those skilled in the art are familiar with the use of various coding and scripting techniques to perform desired operations within a computing system, and a detailed description is unnecessary to an understanding of the invention disclosed herein; all such embodiments, however, are intended to be within the scope of the claims recited hereinafter.

Turning now to FIGURES 3-a and 3-b, with continuing reference to FIGURE 1, illustrated is an exemplary method 300 for automatically replicating

database information according to the principles of the present invention. The method 300 can be implemented as a software-defined process, executable by a general or specific purpose computer to perform the operations, or steps, illustrated. Although illustrated as a series of consecutive steps, those skilled in the art will recognize that the exemplary method 300 merely illustrates the principles of the present invention, and that other alternative, yet substantially equivalent, operations can be used to accomplish an equivalent result.

Initially, in Step 310, it is determined whether there are any custom database information subscriptions. The Database Replication Server 120 can perform this step through a query of the Subscription Database 110. If there are no custom database information subscriptions, the method 300 is terminated. If a custom database information subscription does exist, the custom attributes are retrieved in Step 320. As noted previously, custom attributes can include the identification of one or more Master Subscription Databases 130, and one or more operations to be performed on the identified Master Subscription Databases 130 to create a Custom Information Database 140. For each Master Subscription Database 130 identified by the custom attributes, a connection is established thereto in Step 330. If a Master Subscription Database 130 is physically remote from the Database Replication Server 120, a connection thereto can be established using conventional, or proprietary, network protocols using methods known to those skilled in the art.

In an exemplary embodiment, the method 300 then creates one or more local databases (Step 340) for storing a replica of each of the Master Subscription Databases 130 identified by the custom attributes. Next, in Step 360, it is determined whether the custom attributes specify the need for merging, partitioning, or filtering operations on the identified Master Subscription Databases 130. If not, the method 300 can be terminated, or, in particularly embodiments, the local databases can be published (Step 365). Rather than allowing a user direct access to the local databases, a copy of the local databases can be stored in a location accessible to the user; such location can be on a remote server.

Next, in Step 370, it is determined from the custom attributes whether the custom database information subscription requires a partition operation. If so, a partitioned copy of the identified Master Subscription Database 130 is created (Step 375). Regardless of the need for partitioning, it is also determined whether
5 the custom database information subscription requires a merging operation (Step 380). If so, the identified databases are merged (Step 285); the merged databases can be stored in a local database and/or in a publication database. In addition to merging and partitioning operations, the custom attributes can also identify one or more filtering operations to be performed on a database (Step
10 380). If a filtering operation is identified, the local copy of the identified Master Subscription Database 130 is filtered using the specified criteria (Step 395); the resulting database can be stored in a local database and/or in a publication database.

The exemplary method illustrated in FIGURES 3-a and 3-b is repeated for
15 each custom database information subscription (identified in Subscription Records 115). Once the necessary local databases and/or publication databases are established, conventional methods can be used to maintain any desired synchronization of data stored in the local and/or publication databases with the Master Subscription Databases 130.

20 Finally, reference is made to FIGURE 4, which illustrates a block schematic of an exemplary database replication server 420 and illustrative custom attribute operations 401 and 402 for automatically replicating database information to generate a custom information database 440 according to the principles of the present invention. As previously noted, a Database Replication
25 Server can be a general, or specific, purpose computing system operative to perform the software-definable processes disclosed herein; the architecture and operation of general and specific purpose computers are known to those skilled in the art and, thus, a detailed description herein is not necessary to an understanding of the present invention. In general, as shown in FIGURE 4, a
30 Database Information Server 420 includes a processor 421, volatile memory 422, such as random access memory (RAM), and a non-volatile media 423, such as

a magnetic storage device, for storing a software-definable process, such as the exemplary process illustrated in FIGURES 3-a and 3-b. The processor 421 and volatile memory 422 are operative to retrieve and execute the software-definable process from the non-volatile media 423, and to access a Subscription Database
5 410. Although illustrated as a discrete element separate from Database Replication Server 420, those skilled in the art will recognize that the contents of the Subscription Database can be stored internal to Database Replication Server 420 in the non-volatile media 423. The Subscription Database 410 can be queried by the Database Replication Server 420 to obtain custom attributes
10 associated with each of a plurality of custom database information subscriptions.

As described hereinabove, the custom attributes for each custom database information subscription include: a) the identification of one or more Master Subscription Databases, and b) at least one operation to be performed on the one or more Master Subscription Databases to create a custom
15 information database. The identified operations are performed by the Database Replication Server 420 to generate one or more Custom Information Databases 440.

FIGURE 4 also illustrates two custom attribute operations examples (401, 402). The examples 401 and 402 illustrate how different operations can be
20 combined to generate a Custom Information Database containing desired database information from two Master Subscription Databases. In example 401, a merge operation is performed on Master Subscription Databases 430-A and 430-B to generate an Intermediate Database 435-A/B. A filter operation is then performed on Intermediate Database 435-A/B to extract the desired database
25 information, which is stored in Custom Information Database 440-A/B. The example 402 performs independent filter operations on Master Subscription Databases 430-C and 430-D, the results of which are stored in Intermediate Databases 435-C and 435-D, respectively. A merge operation is then performed on Intermediate Databases 435-C and 435-D to combine the database
30 information into the desired Custom Information Database 440-C/D. Those skilled in the art will recognize that the example operations 401 and 402 are

merely illustrative of the principles of the present invention, and that the possible combinations of operations that can be identified using custom attributes to define a desired Custom Information Database are essentially limitless.

5 The present invention provides significant advantages to the art of information systems, and, more specifically, to the art of database replication.

The principles of automatic merging, partitioning and filtering of replicated databases can provide a user the capability to automatically create and organize database information tailored to the specific user's needs. Furthermore, the automatic partitioning of such replicated databases provides enhanced security and privacy of the information stored in such databases. Of particular advantage, however, are the efficiencies gained through the ability to conduct a single search, or query, of an automatically generated Custom Information Database 140, rather than the need to conduct multiple searches over multiple Master Subscription Databases 130 in order to obtain the desired information.

15 Although the present invention has been described in detail, those skilled in the art will conceive of various changes, substitutions and alterations to the exemplary embodiments described herein without departing from the spirit and scope of the invention in its broadest form. The exemplary embodiments presented herein illustrate the principles of the invention and are not intended to be exhaustive or to limit the invention to the form disclosed; it is intended that the scope of the invention be defined by the claims appended hereto, and their equivalents.

WHAT IS CLAIMED IS:

1. A method for automatically replicating database information, said method comprising the steps of:

5 querying a subscription database to obtain custom attributes associated with each of a plurality of custom database information subscriptions, said custom attributes for each database information subscription including:

- a) the identification of one or more master subscription databases, and,
- b) at least one operation to be performed on said one or more master subscription databases to create a custom information database, said at
10 least one operation being selectable from the group consisting of:
 - i) merging database information from two or more of said master subscription databases,
 - ii) partitioning database information contained in one or more of said
15 master subscription databases, and
 - iii) filtering database information from one or more of said master subscription databases;

automatically generating a custom information database for each of said custom database information subscriptions by performing said at least one
20 operation on said database information stored in said one or more master subscription databases.

2. The method recited in Claim 1, wherein said merging, partitioning and filtering operations are definable on a record or record field basis.

3. The method recited in Claim 1, wherein said step of generating said
25 custom information database associated with each custom database information subscription comprises the steps of:

creating a local database for storing said custom database information;
establishing a connection to each of said master subscription databases identified by the custom attributes;

retrieving and processing master database records from each of said master subscription databases to yield custom database records according to said at least one operation identified by said custom attributes; and
storing said custom database records in said local database.

- 5 4. The method recited in Claim 3, wherein said step of generating further comprises the step of publishing said custom database information in a remote database.
5. The method recited in Claim 1, wherein said step of generating said custom information database comprises the steps of:
- 10 creating a partitioned database containing a subset of the database records stored in a first one of said master subscription databases; and
 merging said partitioned database with a second one of said master subscription databases to yield said custom information database.
6. The method recited in Claim 1, wherein said step of generating said
- 15 custom information database comprises the steps of:
- creating a merged database containing database records stored in first and second ones of said master subscription databases; and
 filtering selected records from said merged database to yield said custom information database.
- 20 7. The method recited in Claim 1, wherein said step of generating said custom information database comprises the steps of:
- filtering selected records from a first one of said master subscription databases;
- filtering selected records from a second one of said master subscription
- 25 databases; and
 merging said selected records filtered from said first and second master subscription databases to yield said custom information database.

8. The method recited in Claim 1, wherein said master subscription databases comprise records including data selected from the group consisting of:

- 5 telephone information;
 address information; and
 historical account information.

9. A database replication server for automatically replicating database information, said database information server comprising:

- 10 a processor;
 volatile memory; and
 a non-volatile media for storing a software-definable process, said processor and volatile memory being operative to retrieve and execute said software-definable process to perform the acts of:

- 15 querying a subscription database to obtain custom attributes associated with each of a plurality of custom database information subscriptions, said custom attributes for each custom database information subscription including:

- 20 a) the identification of one or more master subscription databases, and,
 b) at least one operation to be performed on said one or more master subscription databases to create a custom information database, said at least one operation being selectable from the group consisting of:

- 25 i) merging database information from two or more of said master subscription databases,
 ii) partitioning database information contained in one or more of said master subscription databases, and
 iii) filtering database information from one or more of said master subscription databases; and

- 30 automatically generating a custom information database for each of said custom database information subscriptions by performing said at least one

operation on said database information stored in said one or more master subscription databases.

10. The database replication server recited in Claim 9, wherein said merging, partitioning and filtering operations are definable on a record or record field basis
- 5 11. The database replication server recited in Claim 9, wherein said act of generating said custom information database associated with each custom database information subscription comprises the acts of:
- creating a local database for storing said custom database information;
 - establishing a connection to each of said master subscription databases
 - 10 identified by the custom directory attributes;
 - retrieving and processing master database records from each of said master subscription databases to yield custom database records according to said at least one operation identified by said custom attributes; and
 - storing said custom database records in said local database.
- 15 12. The database replication server recited in Claim 11, wherein said act of generating further comprises the step of publishing said custom database information in a remote database.
13. The database replication server recited in Claim 9, wherein said act of generating said custom information database comprises the acts of:
- 20 creating a partitioned database containing a subset of the database records stored in a first one of said master subscription databases; and
- merging said partitioned database with a second one of said master subscription databases to yield said custom information database.
14. The database replication server recited in Claim 11, wherein said act of
- 25 generating said custom information database comprises the acts of:

creating a merged database containing database records stored in first and second ones of said master subscription databases; and

filtering selected records from said merged database to yield said custom information database.

- 5 15. The database replication server recited in Claim 9, wherein said act of generating said custom information database comprises the acts of:

filtering selected records from a first one of said master subscription databases;

- 10 filtering selected records from a second one of said master subscription databases; and

merging said selected records filtered from said first and second master subscription databases to yield said custom information database.

- 15 16. The database replication server recited in Claim 9, wherein said master subscription databases comprise records including data selected from the group consisting of:

telephone information;

address information; and

historical account information.

FIG. 1

100

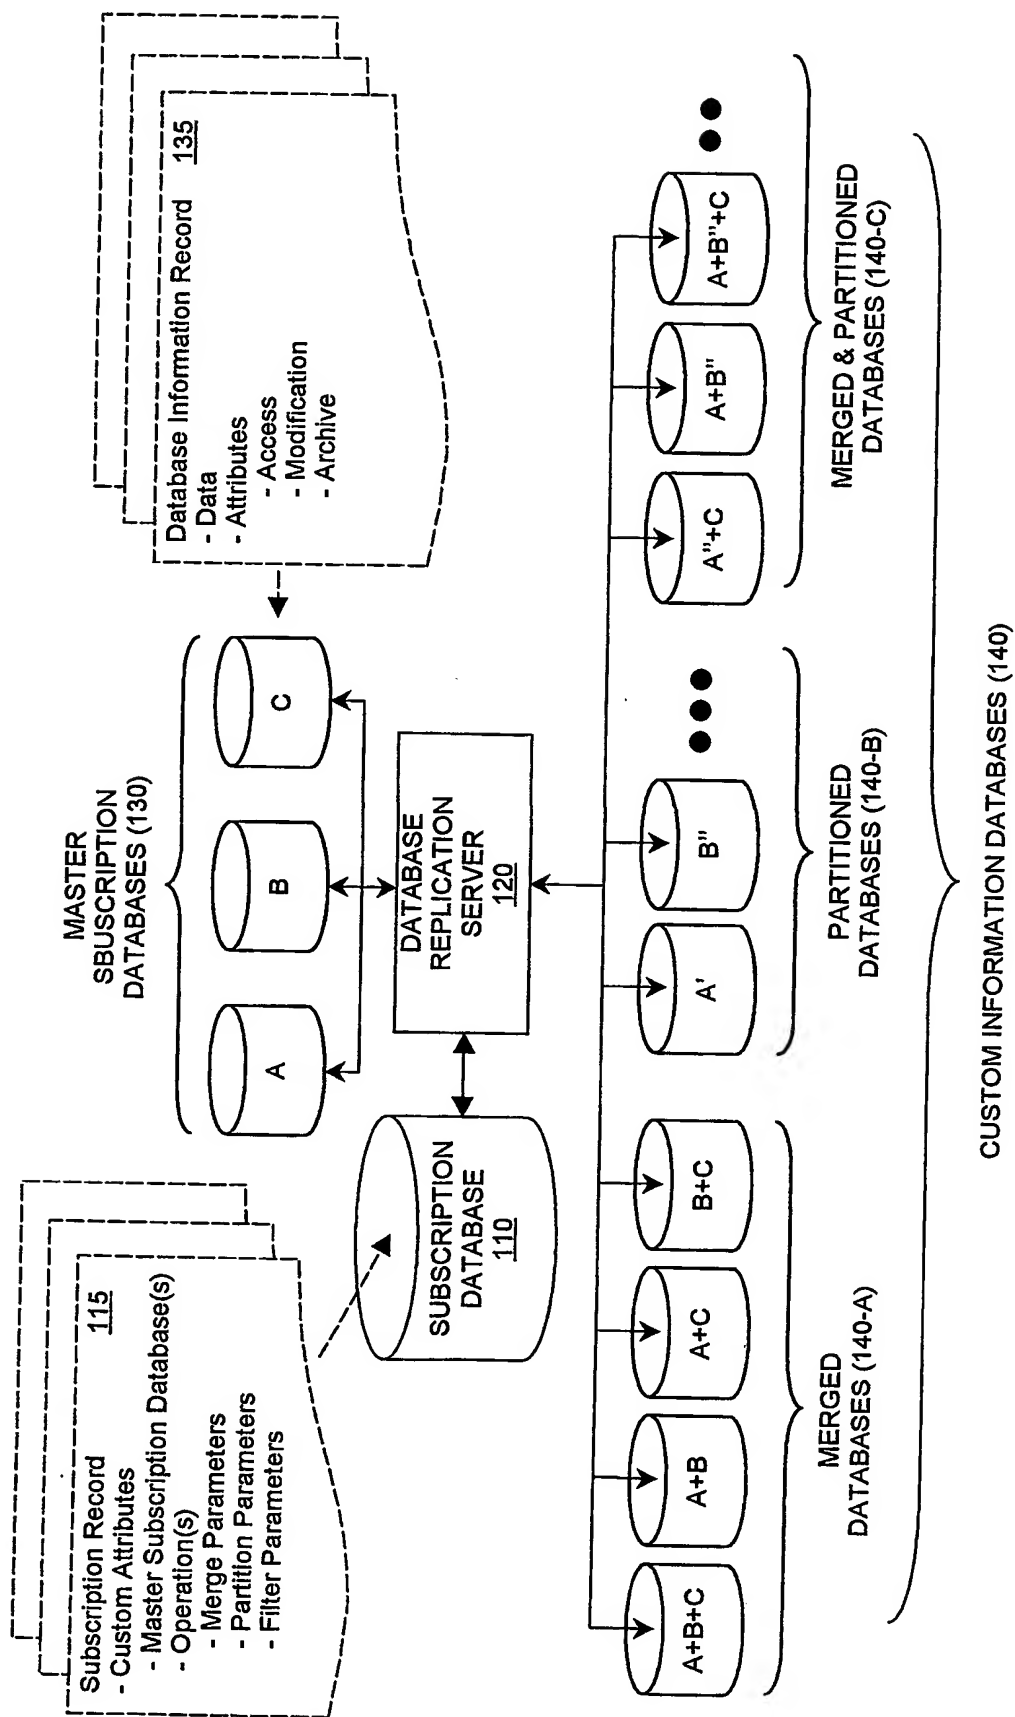
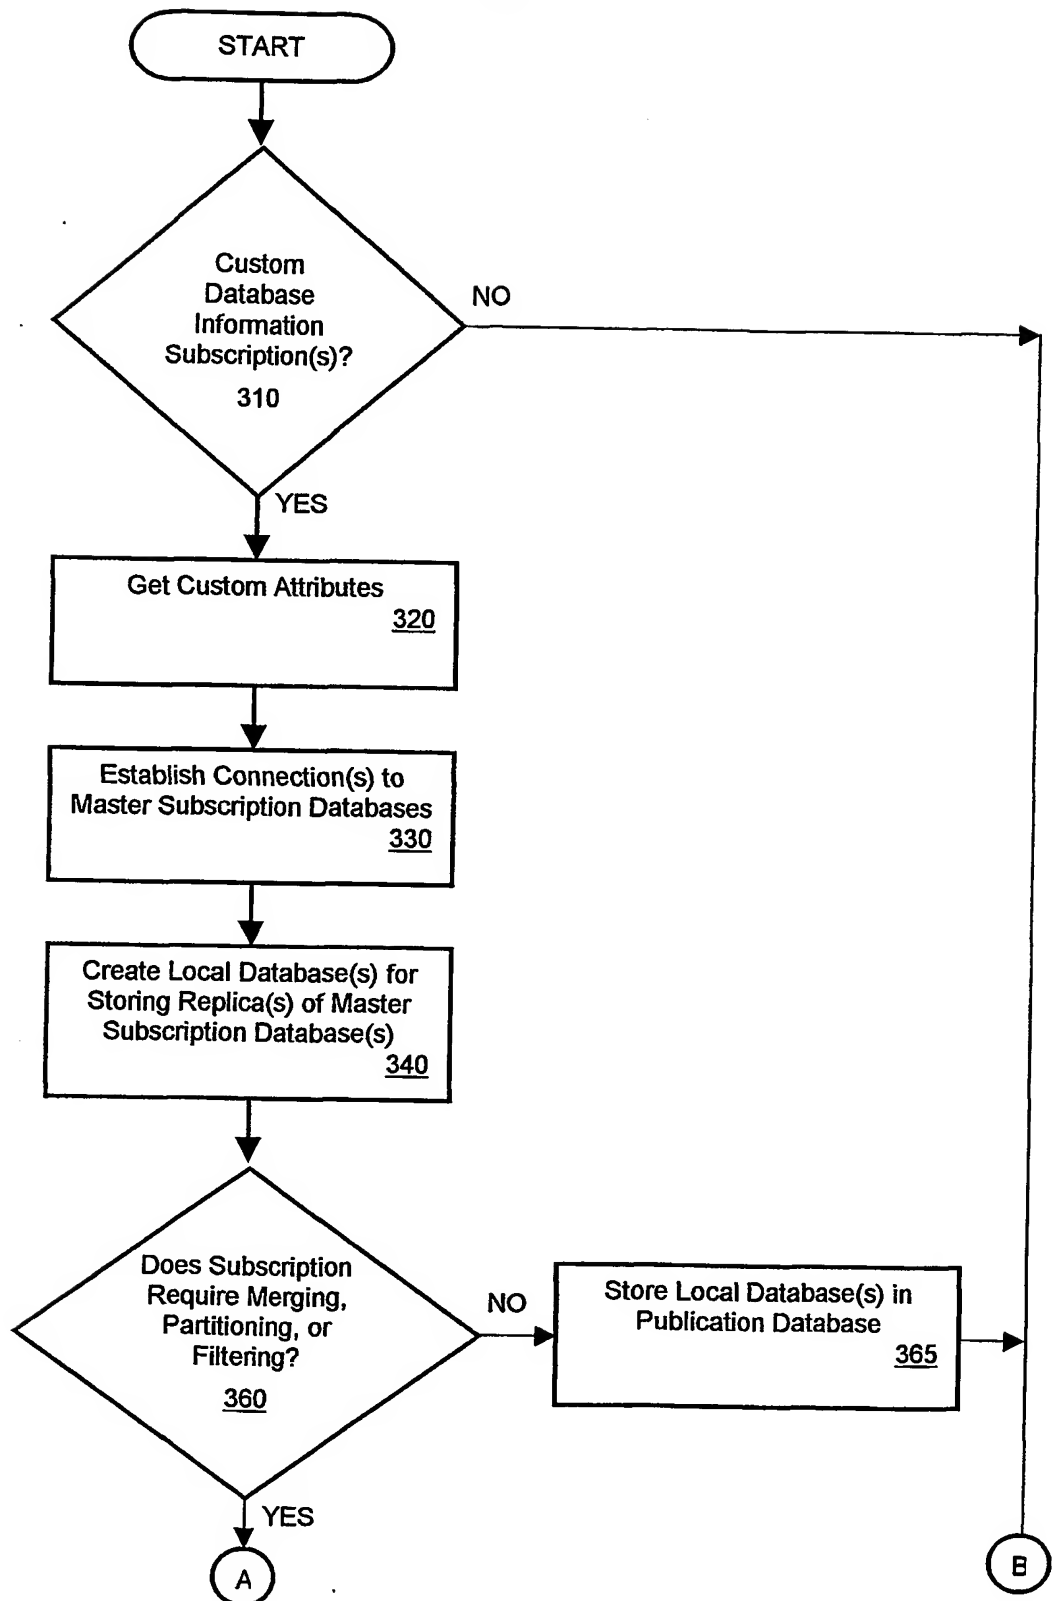


FIG. 3-a300

4/5

FIG. 3-b